Volume Title: Training and the Private Sector

Volume Author/Editor: Lisa M. Lynch

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-49810-7

Volume URL: http://www.nber.org/books/lync94-1

Conference Date: December 15-17, 1991

Publication Date: January 1994

Chapter Title: Productivity Changes without Formal Training

Chapter Author: Andrew Weiss

Chapter URL: http://www.nber.org/chapters/c8780

Chapter pages in book: (p. 149 - 160)

5 Productivity Changes without Formal Training

Andrew Weiss

To check the effects of formal job training programs, it is useful to have a benchmark—a measure of productivity changes that would have occurred without formal training. This study provides such a benchmark by reporting changes in the productivity of workers who do not participate in formal training programs. Those workers may be benefiting from informal training: learning-by-doing or learning by watching other people doing the same or similar tasks.

Standard measures of the return to formal training may be overstating the true return in two ways. First, if formal training is a substitute for informal training, then the forgone benefits from informal training should be subtracted from the returns to formal training. Second, if informal training is bundled with formal training programs and if that informal training would have occurred in the absence of formal training, then standard measures of returns to formal training will again be overstated: The productivity gains from the bundled informal training should be subtracted from the estimated returns to formal training.

The effects of informal training on labor productivity may be large. For instance, perhaps the best-documented finding in the industrial engineering literature is that production costs consistently decline by 10–30 percent every time cumulative output doubles.[1] These declines are observed even when researchers correct for the effects of capital investments.[2]

Andrew Weiss is professor of economics at Boston University.

The author is grateful to Lisa Lynch for valuable comments on an earlier draft of this paper and for suggesting the term "informal training" to summarize learning on the job. Earlier versions of this paper were presented at NBER/CEP Conference on International Comparisons of Private Sector Training, at the Centre for Economic Performance, London School of Economics, December 16, 1991, at the National Bureau of Economic Research, and at Dartmouth College. This paper benefited from substantial research help from Alexandra Lomakin.

1. Perhaps the most forceful advocates of the experience curve effect have been members of the Boston Consulting Group (BCG). As early as 1973 the BCG maintained, "It [the experience curve] is so universal that its absence is almost a warning of mismanagement or misunderstanding [how-

The most commonly accepted explanation for these cost reductions is learning-by-doing. If this explanation is correct, then biases in standard estimates of the return to formal training are likely to be large. In this paper we estimate the effects of informal training on the productivity of a sample of newly hired production workers. We find that the effects are large during the workers' first few months on the job but fall rapidly. There are no noticeable improvements in the productivity of workers with more than six months of job tenure. Thus for the workers in our sample, learning-by-doing does not cause the significant gains associated with cumulative output.

Most previous estimates of the effects of learning-by-doing are derived either from aggregate data at the firm, or industry, level or from wage equations which implicitly assume that wages are proportional to productivity. That assumption is invalid if workers acquire firm-specific human capital. If workers have skills that are specific to a firm, equilibrium is generally characterized by wages that systematically deviate from the marginal product of the worker: Experienced workers will be paid less than the full value of their output (they do not capture the entire return to their firm-specific training), while newly hired workers could be paid more or less than the value of their output, depending on how the costs and benefits of firm-specific training are distributed.

Similarly, even in the absence of firm-specific training or human capital, informational asymmetries may lead to wages that are not proportional to the marginal products of workers. Firms may commit to wages for experienced workers that are above the marginal products of those workers as a means of attracting workers who are unlikely either to quit or to be fired. These high wages would be financed by paying newly hired workers wages that are less than those warranted by their expected productivity. Steep wage-tenure profiles also discourage quits, absenteeism, or shirking by those workers.[3] Consequently there are serious difficulties with using wage changes as proxies for changes in labor productivity.

ever] the basic mechanism that produces the experience curve effect is still to be adequately explained" (quoted in Abernathy and Wayne 1974). Economists and managerial consultants have typically assumed that learning-by-doing generates the learning curve and have made recommendations both for corporate and national policy based on that assumption. See Arrow (1962) and Spence (1981) for particularly insightful treatments of the implications of learning-by-doing for corporate strategy.

2. Berndt (1991) has argued that most estimates of learning effects are biased because they are also capturing scale effects. Studies that have attempted to measure scale effects separately from learning effects include Lieberman (1984) and Joskow and Rose (1985).

3. In Guasch and Weiss (1980, 1982) firms pay the workers that they keep wages that exceed the value of their output and pay newly hired workers less than the value of their output. This wage schedule deters applications from workers who believe they will be unsuccessful (or who think they are likely to quit). Similarly, in Salop and Salop (1976) and Salop (1973, 1979) wages rise faster than productivity, either to attract workers with low quit propensities or to reduce the probability of a given worker's quitting. In all those models, experienced workers are paid more than the value of their marginal product. Greenwald (1986) presents a model in which firms pay their experienced workers less than the value of their output. In the Greenwald model, firms gain private

We avoid those problems by using direct measures of physical output. However, these direct measures have their own problems. First, we have data from only three factories; it is possible that the results are idiosyncratic to either the production technologies in those factories or to the specific workers studied. A related problem is that because the productivity of individual workers was carefully monitored, there may have been strong peer-group pressure not to "break the rates," i.e., not to produce at too fast a pace. If measuring output seriously depresses output growth, then using direct measures of output to estimate the effects of learning-by-doing on productivity would underestimate those effects for workers whose output is not directly measured. Peer pressure to work at less than one's optimal pace is likely to have its greatest effects on the most able workers. Indeed, in Weiss (1992) we showed that productivity improvements are negatively correlated with initial productivity levels. This convergence to a standard pace could explain the lack of productivity improvements among the experienced workers in our sample. Although peer pressure to meet a standard is not a special characteristic of the establishments in our study, it is a common feature of labor relations in firms, and thus is likely to affect the relationship between experience and labor productivity in many different settings.

5.1 Evidence from Production Workers

This study shall use data from groups of workers hired at three facilities of a large telecommunications manufacturer. The telecommunications industry had exceptionally rapid rates of growth in both total factor productivity and in labor productivity during the years we studied: Value added per worker grew 80.6 percent between 1977 and 1982, and value added per hour grew 84.2 percent. In the manufacturing sector as a whole, those growth rates were 63 percent and 66.7 percent, respectively. If the growth in labor productivity is due to learning-by-doing by production workers, we would expect to find that experience strongly affects labor productivity for the workers in our sample.

The workers are grouped in the following way: workers hired at plant A in 1977, workers hired at plant A in 1979, workers hired at plant B, and workers hired at plant C. Plant A assembled components, plant B assembled small final products, and plant C assembled large computers. These data were first analyzed in 1980, in an effort to improve the hiring criteria of the firm. (At the time the only objective criterion for hiring workers was scores on a simple

information about their employees, generating a "winner's curse" situation which deters other firms from bidding for those workers. Consequently, all firms can underpay their experienced workers. Firms offer inexperienced workers wages that exceed their expected productivities, by an amount sufficient to offset (in expected present value) the future profits the firms will get because of the private information they will acquire about those workers.

Table 5.1 **Characteristics of the Sample**

Characteristic	Mean	Standard Deviation
Mean age (years)	25.00	7.3
Mean education (years)	12.10	1.2
Fraction male	0.43	
Fraction married	0.44	
Fraction employed at application	0.51	
Median pay increase (%)	103	

Note: Sample consists of usable output data for approximately 2,000 workers, the exact number depends on the independent variable being considered.

dexterity test.)[4] The data were originally collected to help administer the wage incentive plan. Relevant characteristics of the sample are presented in table 5.1.

These data have several significant virtues. We have direct measures of the output of those semiskilled production workers in the sample who were assigned to "bench" as opposed to assembly line jobs. Consequently the workers in our subsample had considerable control over the pace at which they worked. These workers were paid piece rate until they had worked for a calendar month in which their output was 82.5 percent of the expected output for an experienced worker on that job. At that point the worker was assigned to a pay group and paid according to the output of her group. The average pay group had 125 members, so that once a worker joined a pay group her performance had only a very small effect on her output. Almost all the workers were in pay groups by their third or fourth month on the job. The lack of direct financial incentives for these workers is typical of U.S. production workers. It is unusual for U.S. workers to receive individual incentive pay or to have their pay directly linked to their output.

These workers were *atypical* in that their output was carefully monitored: the output of each worker in a pay group has an effect, albeit small, on the pay of every worker in that group. Because of the direct connection between performance and pay, great care was taken to determine the expected level of output for an experienced worker on each job. These standards are set by the industrial engineering staff at each location. Engineers from the firm's headquarters review the standards. The typical revision of a job standard involves a correction of less than 2.5 percent. Standards that workers perceive as unusually demanding can be appealed through a union grievance procedure. More attention is paid to achieving uniformity of rates within a plant than uniformity

4. In plants A and B, a two-part dexterity test was administered. The first part was a pins test in which applicants were asked to insert small pins into sleeves. Their score was the number of pins they were able to insert during a three-minute period. The second part of the test measured how many screws an applicant was able to fully insert into threaded sleeves during a three-minute period. In plant C, only the screws test was administered.

across plants. It is difficult for workers or officials of the union local to compare rates for jobs that are 1,000 miles apart. It is easier to see that the person sitting next to you is working at a significantly slower pace. Thus there is more pressure by workers on the industrial engineers to achieve rate uniformity within a plant than across plants. Second, the differences among jobs are greater across plants than within plants. These differences increase the relative difficulty of achieving uniformity of rates across plants.

For workers assigned to a new job, the expected output for each worker is adjusted according to the "learning curve" associated with that job. Industrial engineers estimate the job-specific learning curve by determining the proportion of the expected output of an experienced worker they expect a newly hired worker to produce when first assigned to the job. The median expected learning times for jobs at plants A, B, and C in this study were 12, 7, and 15 weeks, respectively. Within the entire subsample the expected learning time ranged from 1 to 36 weeks.

5.2 Effects of Learning-by-Doing

The data in table 5.2 describe median changes in the physical output of individual workers. (Note, we are not reporting the change in median productivity. Those data are reported in table 5.4, below.) Thus the data in table 5.2 remove most of the effects of changes in labor-force composition on changes in average productivity. Sample selection bias remains only if employment or hiring decisions are correlated with *future changes* in productivity. Furthermore, that correlation must not be a general feature of labor markets. One source of sample selection bias that is likely to have remained in the data is differential quit rates. Quit rates are likely to be higher among workers who anticipate having the most difficulty meeting performance standards. These will tend to be workers whose initial productivity was low and who anticipate relatively small increases in their future productivity. In general, the high quit rates of workers who anticipate not achieving sufficient growth in productivity to meet the standards will lead to overestimates of the effects of experience on productivity. These effects could be offset, or even reversed, if the workers whose initial performance was high had lower rates of growth of output. In-

Table 5.2 **Median Percentage Change in Productivity**

Tenure (months)	Plant A, 1977	Plant A, 1979	Plant B	Plant C
1–2	10.8	18.5	45.2	35.4
2–3	2.0	2.0	13.0	8.3
3–4	0.8	0.8	6.4	3.2
4–5	0.4	0.4	3.0	1.9
5–6	0.0	0.4	1.0	0.3
6 7	—	—	—	0.1

deed, we find that relationship in our data (see Weiss 1992), thus the direction of the bias from the correlation between quit rates and expected growth in productivity is unclear.[5]

Plant A had data on the monthly output during the first six months of employment both for workers hired in 1977 and for workers hired in 1979. For workers who were hired in 1977 at plant A and who remained with the firm for six months, the median change in their output between their fifth and sixth months on the job was precisely zero. The median change in output from their sixth month on the job to the last period for which we had records was −1 percent. (This last period varied across workers, depending on when in 1977 they were hired.)

For workers hired in plant A in 1979, the median change in productivity between the fifth and sixth months was 0.4 percent. This contrasts with a median increase in productivity of 18 percent from the first to second months of employment. Because of an unanticipated fall in product demand, few of those workers stayed with the firm beyond six months. But extrapolating from trend, we would expect trivial increases in output past the sixth month on the job. The sample sizes for these cohorts varied somewhat during the sample period due to quits or to workers being reassigned to assembly-line jobs or to jobs at which individual output could not be directly measured. To give some idea of the sample sizes: in the 1977 sample we had usable output data for 585 workers during their sixth month on the job. For the 1979 sample we had usable output data for 308 workers during their sixth month on the job. (Those workers were on bench jobs, so their output could be measured, and were on the same job for the entire month.)

Workers at plant B had much more rapid increases in their productivity. The management at this plant had a reputation for pressing workers to work exceptionally fast. From the first to the second month of employment the median change in output was 45 percent. This relatively high increase in output occurred despite this plant's having the simplest jobs of any in our sample. If we had used only the learning curves to predict changes in productivity, we would have expected growth rates in output to be the *lowest* at plant B. Instead they were the *highest*. However, by the sixth month, productivity growth had fallen to 1 percent. Extrapolating from trend we would expect almost no growth in productivity after the sixth month on the job. Unfortunately we did not have data on individual output after the sixth month for workers in this

5. This observed relationship could exist because the effect of expected productivity gains on quit behavior is greater for workers with initially low levels of performance. In the extreme case, the quit propensities of workers whose productivities were well above the standard would not be affected by differences in anticipated increases in productivity, while among workers whose initial performance was low, only those who anticipated large gains in their productivity would stay. These differences in quit propensity could explain why productivity growth is negatively correlated with initial productivity.

plant. As in plant A, the sample size varied with tenure at the plant. For instance, we had usable output data for 182 workers during their fourth month on the job.

The demands on the workers in plant B were also reflected in a relatively high quit rate. The quit rate at this plant was 22 percent during a worker's first six months of employment. At plants A and C the quit rates were 9 and 12 percent, respectively. The high quit rate at plant B might give rise to large overestimates of the effect of experience on labor productivity for a randomly selected worker. Because of the difficult standards at plant B, an exceptionally large proportion of the workers at that plant may be quitting because they do not expect to achieve sufficient productivity gains to meet the plant's standards.

Workers in plant C had the most sophisticated jobs of any workers in our sample. In that plant the median increase in output was 35 percent from a worker's first to second month on the job. By the sixth month, productivity changes had effectively ceased. At this plant we had usable output data for 178 workers during their sixth month on the job.

5.3 Changes in Work-Force Composition

In the previous section we restricted our attention to changes in the productivity of workers. However, cumulative output may provide information to workers and firms which enables them to better sort themselves. One of the benefits of production is that it can improve the match between workers and jobs. The least productive workers may be pressed to work at a pace that they find so difficult to maintain that they are induced to quit.

In table 5.3 we computed probit estimates of the probability of a worker's quitting within her first six months on the job.[6] The quit function with respect to first-month output has an interior minimum. While the best and the worst workers are most likely to quit, the workers who are least likely to quit are more productive than the average workers. The workers with the lowest expected quit probabilities had productivity levels, during their first month of employment, that were roughly one standard deviation above the mean productivity for the entire sample. Most quits are among workers with relatively low initial rates of output. Thus, for the firm we studied, quits by production workers appear to improve the quality of its work force.

The correlation between quit probabilities and some of the other independent variables is of independent interest. The strongest finding is that workers who quit a job in order to take this job are less likely to quit than are workers who were unemployed at the time they applied for this job. The obvious expla-

6. In estimating the quit equations, we excluded from the sample any workers laid off before they completed six months of employment with the firm. We also excluded workers that would have been laid off before completing six months of employment with the firm, had they not quit first.

Table 5.3 Probit Estimates for Quits within First Six Months on the Job (standard errors in parentheses)

Independent Variable	Estimate
First-month output[a]	−0.335 (0.164)
(First-month output)^2	0.123 (0.075)
Total years of schooling	−0.022 (0.017)
Postsecondary education	0.001 (0.004)
Employed at application[b]	−0.079 (0.029)
Male	0.034 (0.032)
Married	−0.028 (0.028)
Age	−0.002 (0.0164)
Score on dexterity test	0.003 (0.003)
Job complexity[c]	0.023 (0.022)
Plant A dummy variable	0.024 (0.048)
Plant B dummy variable	0.098 (0.054)
Total years of schooling interacted with job complexity[d]	−0.001 (0.002)

[a]A measure of the worker's output during her first month of employment, adjusted by the industrial engineers for the difficulty of learning the job.

[b]A dummy variable taking the value 1 if the worker was employed at the time she applied for the job.

[c]The logarithm of the number of weeks the industrial engineering staff estimates that it takes to be fully trained on the worker's job.

[d]Difference between the worker's schooling and mean schooling × difference between the complexity of the job and mean job complexity.

nation for this result is that the worker-job match for workers who were employed when they applied for these jobs is likely to be better than for workers who were unemployed. Employed workers only quit their jobs if the anticipated increase in utility outweighed their costs of changing jobs, thus those workers are likely to have superior job matches and lower costs of quitting. For determining quits in our sample, the matching effect appears to outweigh differences in the costs of changing jobs. We also find that men are more likely to quit than are women; this difference in quit propensities may be due either to better alternative opportunities for men or to the discomfort felt by men in working in a predominantly female work force: While the new hires were 44 percent male, the existing labor force was less than 20 percent male. Finally, job complexity does not seem to reduce quits, nor does matching the better-educated workers to the more complex jobs appear to have a significant effect on quits.

5.4 Combined Sorting and Learning Effects

Table 5.4 describes the differences in the output of various tenure groups. No attempt is made to correct for changes in labor-force composition: The data illustrate the combined sorting and learning effects of experience on labor productivity. For instance, we can see that the expected output of a randomly selected worker with three months of tenure at this establishment is 11.7 percent higher than that of a randomly selected worker with two months of tenure. Workers with six months of tenure are, on average, 1.7 percent more productive than workers with only five months of tenure. Workers with more than six months of experience are not any more productive than workers with six months of experience. Note that when we are comparing groups of workers with different amounts of tenure, the composition of the group has in general changed over time. Some of the workers quit or were laid off during the first six months they were employed. Those workers will be included in the computation of median and mean productivities for the months before they quit or were laid off but not for the month they were laid off (or for any later month). Consequently, the data in table 5.4 include the effects both of changes in the composition of the work force and of changes in the performance of individual workers. (The data in table 5.4 describe changes in the means and medians rather than the median of the changes as in table 5.2.)

During their first month of employment, almost all the workers spent some of their time watching other workers. Therefore, measuring the changes in output from month one to month two for many of the workers involves a comparison of their output during their first one or two weeks on the job with their output during their next four weeks on the job. Since we do not know how much time workers spent watching other workers doing their jobs, versus doing the job themselves, it is difficult to precisely measure the effect of experience on the change in output. For the workers in our sample, this problem is unlikely to be of great importance after their first month on the job.

Table 5.4 Changes in Productivity with Tenure Combined Sorting and Learning Effects

Period over Which Change is Measured	Change in Median Hourly Output (%)	Change in Mean Hourly Output (%)
Month 1 to 2[a]	23.2	43.9
Month 2 to 3	5.0	11.7
Month 3 to 4	2.1	6.0
Month 4 to 5	1.0	3.5
Month 5 to 6	0.27	1.7
Month 6 to 18 (plant A)	−1.01	−0.002
Month 6 to 7 (plant C)	0.1	−0.6
Month 7 to 8 (plant C)	−0.6	−0.7
Month 5 to 6 (plant A; workers in lowest quartile in first month)	1.1	5.0

[a]Includes only workers in jobs requiring less than four weeks of training.

While the expected productivity of a labor force with two months of job experience is 44 percent greater than that of a labor force with one month of experience, these productivity improvements decline rapidly as we compare more experienced cohorts. Among workers who have been employed for at least six months, we do not see any further increases in productivity.

As discussed above, we are measuring the combined effect both of changes in the performance of individual workers and of changes in the composition of the work force. Because the least able workers are most likely to quit, these data are likely to give an upwardly biased estimate of the effect of learning-by-doing on productivity during a worker's first six months of employment. In other words, even if none of the workers increased individual productivity in any period, we would find increases in productivity in each period, because the least able workers were more likely to quit in each period.

It is relatively easy to compare the data in table 5.4 with data from experience-curve studies which also do not adjust for differences in quit propensities among more and less able workers. The data in table 5.4 suggest that neither favorable selection nor learning-by-doing of production workers can explain the experience curve. In the data we analyzed, there was no evidence of a correlation between the cumulative output of the product associated with a worker's job and labor productivity on that job. There is, however, one group whose measured productivity shows significant growth even after five months on the job. These are workers whose initial productivity placed them in the lowest quartile (adjusting for the difficulty of learning their job). The output of those workers increases by 5 percent from their fifth to sixth month on the job. For data on low-productivity workers, we only included workers in plant A. In plant B the productivity of the bottom quartile was so low that workers who were still with the firm at the end of six months would have to have been increasing their output significantly to avoid pressure to quit. For plant C the small size of the sample and the long average learning curves precluded doing any meaningful analysis of the performance of workers with low initial productivity. Because those workers are likely to be ones for whom the learning curve was underestimated, in plant C this measurement error could also distort the change in output measures from their fifth to sixth month of employment. Only plant C had jobs for which the expected learning period exceeded six months.

5.5 Concluding Remarks

The main conclusions we would draw from this study are that rapid productivity growth occurs during the first month that production workers are employed, even without formal job training programs. However, productivity growth falls rapidly and effectively stops by the sixth month on the job. The fall may be due to peer pressure not to "break the rates."

The data we have presented also cast some light on the extent to which aggregate productivity growth is due to learning-by-doing. It appears that learning-by-doing by production workers has only a small effect on productivity growth. Despite the very rapid rates of productivity growth in the industry we studied, there was almost no net change in the output of workers after they gained four to six months of experience on their jobs. If the widely observed correlation between cumulative output and labor productivity is due to factors other than learning-by-doing, then policies designed to increase market share (either at the firm or national level) as a means of reducing future costs may be misguided. Policies that directly address factors, such as cumulative engineering inputs, that are causing the correlation between cumulative output and labor productivity are likely to be more effective.

References

Abernathy, William J., and Kenneth Wayne. 1974. Limits of the learning curve. *Harvard Business Review* 52, no. 5 (September/October): 109–19.

Arrow, Kenneth. 1962. The economic implications of learning by doing. *Review of Economic Studies* 29 (June): 115–73.

Berndt, Ernst. 1991. The practice of econometrics classic and contemporary. Reading, Mass.: Addison Wesley.

Greenwald, Bruce. 1986. Adverse selection in the labor market. *Review of Economic Studies* 53:325–47.

Guasch, J. Luis, and Andrew Weiss. 1980. Wages as sorting mechanisms in competitive markets with asymmetric information. *Quarterly Journal of Economics* 94 (3): 453–66.

———. 1982. An equilibrium analysis of wage-productivity gaps. *Review of Economic Studies* 49:484–97.

Joskow, Paul L., and Nancy Rose. 1985. The effects of technological change, experience and environmental regulation on the construction cost of coal burning generating units. *Rand Journal of Economics* 16, no. 1 (Spring): 1–27.

Lieberman, Marvin. 1984. The learning curve and pricing in the chemical processing industries. *Rand Journal of Economics* 15, no. 2 (Summer): 213–28.

Salop, Joanne, and Stephen Salop. 1976. Self selection and turnover in the labor market. *Quarterly Journal of Economics* 90, no. 4 (November): 619–27.

Salop, Stephen. 1973. Wage differentials in a dynamic theory of the firm. *Journal of Economic Theory* 6 (4): 321–44.

———. 1979. A model of the natural rate of unemployment. *American Economic Review* 69 (1): 117–25.

Spence, A. Michael. 1981. The learning curve and competition. *Bell Journal of Economics* 12, no. 1 (Spring): 49–70.

Weiss, Andrew. 1992. Productivity changes without formal training. Working Paper. The Ruth Pollak Series in Economics. Working Paper no. 9. Boston, Mass.: Boston University.